



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 27 May 2003



Patents Act 1977
Rule 16

THE PATENT OFFICE
G
- 8 APR 2003
NEWPORT

The
Patent
Office

1/77

08APR03 E798502-1 D00611
P01/7700 0.00-0308035.5

Request for grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1. Your reference	GB920020070GB1		
2. Patent application number <i>(The Patent Office will fill in this part)</i>	0308035.5		
	E8 APR 2003		
3. Full name, address and postcode of the or of each applicant <i>(underline all surnames)</i>	INTERNATIONAL BUSINESS MACHINES CORPORATION Armonk New York 10504 United States of America		
Patents ADP number <i>(if you know it)</i>	519637001		
If the applicant is a corporate body, give the country/state of its incorporation	State of New York United States of America		
4. Title of the invention	LIVENESS MONITORING IN A PUBLISH/SUBSCRIBE MESSAGING SYSTEM		
5. Name of your agent <i>(if you have one)</i>	D P Litherland		
"Address for Service" in the United Kingdom to which all correspondance should be sent <i>(including the postcode)</i>	IBM United Kingdom Limited Intellectual Property Department Hursley Park Winchester Hampshire SO21 2JN		
Patents ADP number <i>(if you know it)</i>	6603898002		
6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and <i>(if you know it)</i> the or each application number	Country	Priority App No <i>(if you know it)</i>	Date of filing <i>(day/month/year)</i>
7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	No of earlier application	Date of filing <i>(day/month/year)</i>	

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:
 a) any applicant named in part 3 is not an inventor, or
 b) there is an inventor who is not named as an applicant, or
 c) any named applicant is a corporate body.)

Yes

9. Enter the number of sheets for any of the following items you are filing with this form.
 Do not count copies of the same document

Continuation sheets of this form

Description 12

Claim(s) 6

Abstract 1

Drawing(s) 5 + 5

11

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 6

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents
(please specify)


I/We request the grant of a patent on the basis of this application


 Signature

D P LITHERLAND

7 April 2003
Date

12. Name and daytime telephone number of person to contact in the United Kingdom J N Watson
01962 818955



The
Patent
Office

7/77

Patents Act 1977
Rule 15

Statement of inventorship and of right to grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1. Your reference	GB920020070GB1
2. Patent application number <i>(if you know it)</i>	0308035 5 F 8 APR 2003
3. Full name of the or of each applicant	INTERNATIONAL BUSINESS MACHINES CORPORATION
4. Title of invention	LIVENESS MONITORING IN A PUBLISH/SUBSCRIBE MESSAGING SYSTEM
5. State how the applicant(s) derived the right from the inventor(s) to be granted a patent	By employment and by agreement
6. How many, if any, additional Patents Forms 7/77 are attached to this form?	1
7.	I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.

Signature D P LITHERLAND

7 April 2003
Date

8. Name and daytime telephone number of person to contact in the United Kingdom	J N Watson Tel: 01962 818955
--	---------------------------------

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Boaz CARMELI
(Resident of Israel)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (*if known*)

8605826001

John Justin DUIGENAN
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (*if known*)

8605834001

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

Gidon GERSHINSKY
(Resident of Israel)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (*if known*)

8605842001

REMINDER

Have you signed the form?

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Stephen James TODD
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (*if known*)

8530537001

Graham Derek WALLIS
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (*if known*)

8520132001

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

REMINDER

Have you signed the form?

Patents ADP number (*if known*)

LIVENESS MONITORING IN A PUBLISH/SUBSCRIBE MESSAGING SYSTEM**Field of the Invention**

5 This invention relates to brokered multicast publish/subscribe messaging systems.

Background of the Invention

10 Publish and Subscribe is an effective way of disseminating information to multiple users. Publish/Subscribe applications can help to enormously simplify the task of getting business messages and transactions to a wide, dynamic and potentially large audience in a timely manner.

15 In a publish/subscribe messaging system subscribers register their interest in one or more topics. The broker performs a match of publications to interested subscribers and sends a copy of each publication to the appropriate subscribers. The stream of publication messages is divided into a sequence of packets of sizes that are optimal for the transmission medium being used. To maximise the efficiency of the network utilisation in such a 20 publish/subscribe system it is preferable to multicast the packets that contain the messages which are to be sent to a number of subscribers. Where there is a large number of subscribers for a given topic the network efficiency gain provided by multicast is greater. The broker performs the 25 role of multicast transmitter and the subscribers each perform the role of multicast receiver.

30 In a reliable multicast publish/subscribe system, subscribers request retransmission of any packet that is not delivered. They do this by detecting gaps in the delivery sequence. When a subscriber detects a missing packet it requests retransmission by sending a "negative acknowledgement" or NACK. To avoid the generation of a storm of NACKs when 35 a packet goes missing, the subscribers can use a NACK suppression mechanism, which operates by each subscriber setting a random back-off timer and sending a multicast NACK packet on expiry of the timer. If a subscriber sees another subscriber's NACK packet before its own timer expires, it cancels the timer.

40 However, this approach has the disadvantage(s) that the only feedback that the broker has is the receipt of NACK packets when one or more subscribers fail to receive a packet and the notification during orderly subscriber termination that a subscriber no longer wishes to receive publications matching a particular set of topics. The broker has no guarantee that either of these forms of feedback will be received; no

5

packets may be being dropped and subscribers could fail or disconnect unintentionally. Accordingly, the broker has no knowledge of the current status of the subscribers and is therefore obliged to keep multicasting publications even when no subscribers are actually running, thus reducing the efficiency of such a system.

10

A need therefore exists for efficient liveness monitoring in a reliable multicast system wherein the abovementioned disadvantage(s) may be alleviated.

15

Statement of Invention

20

In accordance with a first aspect of the present invention there is provided a publish/subscribe messaging system, comprising: at least one broker and at least one subscriber, the broker having means for sending a status request message to the subscriber, and means for receiving an indication of liveness of the subscriber.

25

Preferably the subscriber has means for sending a status response message to the broker to indicate liveness.

25

In one embodiment, the means for receiving comprises means for listening on a multicast channel and for determining an indication of non-liveness from failure to receive a response from the subscriber.

30

For example, the broker may listen on the multicast channel and may hear a subscriber "claiming" that it will respond to the broker without any explicit response to the broker being necessary (see later).

35

Preferably the means for sending a status response message to the broker comprises means for suppressing sending of the status response message if at least another subscriber sends a status response message.

40

In a preferred embodiment, the means for suppressing sending of the status response message comprises: means for setting a timer upon receipt of a status request message from the broker; means for sending, on expiry of the timer, a multicast message claiming response to the broker; means for cancelling the timer and discarding the status request message if the subscriber receives a message claiming response from another subscriber; means for sending the status response message to the broker following sending a message claiming response.

In one embodiment, the means for suppressing sending of the status response message comprises: means for setting a timer upon receipt of a

status request message from the broker; means for sending, on expiry of the timer, a status response message to the broker; and means for cancelling the timer and discarding the status request message if the subscriber sees a status response message from another subscriber.

5

Preferably, the broker further comprises means for re-sending the status request message if it does not receive a response thereto.

In one embodiment, the means for suppressing sending of the status response message is arranged to suppress sending of the status response message if at least a desired plurality of other subscribers send a status response message.

In this embodiment, the means for suppressing sending of the status response message may comprise: means for setting a timer upon receipt of a status request message from the broker, the status request message containing a parameter representative of the desired plurality of other subscribers; means for sending, on expiry of the timer, a multicast message claiming response to the broker; means for cancelling the timer and discarding the status request message if the subscriber receives messages claiming response from the desired plurality of other subscribers; means for sending the status response message to the broker following sending a message claiming response.

25

In one embodiment the timer has a random duration.

30

In one embodiment at least one of the subscribers is arranged to maintain an active connection to the broker established during registration, and to use the active connection to indicate liveness to the broker.

35

40

In one embodiment the means for suppressing sending of the status response message comprises: means for checking, upon receipt of a status request message from the broker, whether the subscriber has an active connection to the broker and if so performing one of A) and B): A) sending a multicast response claim message, and sending a status response message to the broker via the active connection, and B) setting a timer and then sending a multicast response claim and a status response message to the broker via the active connection; and means for, following sending of a multicast response message, establishing an active connection to the broker if not already established and sending the status response message to the broker via the active connection.

5 In one embodiment the means for suppressing sending of the status response message comprises: means for checking, upon receipt of a status request message from the broker, whether the subscriber has an active connection to the broker and if so performing one of A) and B): A) sending a status response message to the broker via the active connection, and B) setting a timer and then sending a status response message to the broker via the active connection; and means for establishing an active connection to the broker if not already established and sending the status response message to the broker via the active connection.

10

In one embodiment, the broker is arranged to designate as a primary subscriber the first subscriber to register interest in a topic, and to maintain an active connection to the primary subscriber for sending directly to the primary subscriber a status request message, and in the event of failure of the primary subscriber to send a status request message to at least one other subscriber and to designate as a new primary subscriber the at least one of the other subscribers whose indication of liveness is next first received.

15

In one embodiment, the active connection is a TCP/IP connection.

In one embodiment, the status request message is piggybacked onto another multicast publication message.

20

In one embodiment, the indication of liveness is sent over one of: a UDP connection, and a TCP connection.

25

In one embodiment, the connection over which the indication of liveness is sent is arranged to escalate autonomously from a UDP connection to a TCP connection in the event of no responses being received by the broker within a chosen time period.

30

In accordance with a second aspect of the present invention there is provided a method for liveness monitoring in a publish/subscribe messaging system having at least one broker and at least one subscriber, the method comprising: at the broker, sending a status request message to the subscriber, and receiving an indication of liveness of the subscriber.

35

In accordance with a third aspect of the invention, there is provided a method of liveness monitoring in a publish/subscribe messaging system in accordance with claim 28.

It will of course be appreciated that the invention may be implemented in software.

Brief Description of the Drawing(s)

Embodiments of the present invention will now be described, by way of example only, and with reference to the following drawings:

5

FIG. 1 shows a block schematic diagram of a publish/subscribe messaging system in which embodiments of the present invention may be used;

10 FIG. 2 shows a schematic diagram depicting message flows between components of the system of FIG. 1;

15 FIG. 3 shows a flow diagram depicting method steps of a first technique for liveness monitoring in accordance with an embodiment of the present invention.

15

FIG. 4 shows a flow diagram depicting method steps of a second technique for liveness monitoring in accordance with an embodiment of the present invention; and

20

FIG. 5 shows a flow diagram depicting method steps of a third technique for liveness monitoring in accordance with an embodiment of the present invention.

Description of Preferred Embodiments

25

FIG. 1 shows a brokered publish/subscribe multicast messaging system 100 in which a broker 110 brokers sending of multicast messages from Publisher 1 (publishing information on, for example, the topic of Sport), Publisher 2 (publishing information on, for example, the topic of Stock) and Publisher 3 (publishing information on, for example, the topic of Films & Television) to Subscriber 1 (subscribing to information on, for example, the topics of Sport and Stock), Subscriber 2 (subscribing to information on, for example, the topic of Films & Television) and Subscriber 3 (subscribing to information on, for example, the topic of Sport).

35

As shown in FIG. 2 at 210, Subscriber 1, Subscriber 2 and Subscriber 3 each send a message to the broker 110 to register the respective subscriber with the broker 110, and in response thereto the relevant subscriber receives a message from the broker 110 confirming registration. Thereafter, as shown at 220, each publisher publishes its information to the broker 110, and the broker 110 publishes the information to the relevant subscriber(s) that have registered with the broker to subscribe to such information.

40

As referred to above, if a subscriber detects a missing packet it requests retransmission by sending a "negative acknowledgement" or NACK 230. To avoid the generation of a storm of NACKs when a packet goes missing, the subscribers can use a NACK suppression mechanism, which operates by each subscriber setting a random back-off timer and sending a multicast NACK packet on expiry of the timer. If a subscriber sees another subscriber's NACK packet before its own timer expires, it cancels the timer.

Finally, as shown at 240, Subscriber 1, Subscriber 2 and Subscriber 3 may each send a message to the broker 110 to deregister the respective subscriber from the broker 110, and in response thereto the relevant subscriber receives a message from the broker 110 confirming deregistration.

In the system 100 it is desired, to improve network utilisation and security, to avoid sending multicast packets from the broker when there are no active subscribers. The broker therefore needs to keep track of the set of active subscribers. It is not sufficient to rely on the subscribers un-registering when they are deactivated, because a subscriber may be accidentally disconnected or fail and not get a chance to deregister.

Furthermore, it is important for each subscriber to know if the broker fails and is restarted, so that subscriptions can be re-registered, fresh security keys exchanged and packet sequence numbers can be reset.

The following conditions together preferably indicate the liveness of the system:

Condition 1): *Each subscriber knows that the broker is still active.* This first condition can be evaluated by the in-band receipt of normal data, or periodic "heartbeats" sent by the broker in periods when there is no data to transmit. Each receiver knows what this period is, and if no data or heartbeat arrives from the transmitter within a time related to this period, the receiver takes this as an indication that the transmitter is dead. This is a widely accepted practice, used in channels of IBM's MQSeries ('IBM' and 'MQSeries' are registered trademarks of IBM Corporation) products and need not be described in further detail.

Condition 2): *The broker knows that there is at least one active subscriber.*

To evaluate this second condition the broker needs to periodically receive an indication from at least one subscriber. This can either be prompted by the subscribers or requested by the broker. In the cases where the broker

5 requests status information, the requests can be either be sent in-band or out-of-band. The advantage of sending such requests out-of-band is that this presents less load on the data channel and on the subscriber, which does not need to parse to identify the status request packets within the data stream. Requests can be piggybacked onto other in-band traffic. The disadvantage of sending requests out-of-band is that it does not test the data channel and so does not extend to cover the following third condition.

10 *Condition 3): The broker knows that at least one active subscriber can receive multicast packets.*

15 To evaluate this third condition, the broker needs to send data over the multicast channel, to prove that the multicast channel is operating correctly and that subscribers can receive packets from it, and to receive some feedback from the subscribers.

15 It is therefore desirable to find a cost-effective and scalable means to query subscriber liveness in a reliable multicast system.

20 Following are four techniques that can be used to solve the above problem.

Technique #1

25 When there is data to be sent and there are packet losses, some subscribers will be sending NACK packets. In these conditions the broker can ascertain that there is at least one active subscriber.

30 When there is no data to be sent, or any data transmission is lossless, there will be no NACK packets. It would not be sufficient for the subscribers to use a timeout to trigger the sending of status packets to the broker because this does not prove that the multicast channel is working. The broker therefore needs to send data over the multicast channel, and to receive some feedback from the subscribers. In order to have reliable communication of the feedback, responses can be unicast over a TCP/IP (Transmission Command Protocol / Internet Protocol) connection rather than through the multicast fabric. Alternatively, the responses can be sent using UDP/IP (User Datagram Protocol / Internet Protocol) which is a less reliable point to point protocol. The lower reliability may lead to more requests being generated by the broker; on the other hand, it avoids TCP/IP connection set-up cost. The choice of protocol could therefore be made dependent on the loss rate and number of subscribers and made as a result of dynamic evaluation of these parameters, thereby providing self-optimising characteristics. The broker can escalate from UDP/IP to TCP/IP in the event of no responses being received within an acceptable

time period. It would alternatively be possible in principle to use the reliable multicast protocol to achieve this, but since there is only one intended recipient it is more efficient to use a unicast protocol - hence TCP/IP or UDP/IP.

5

As a result, the broker may periodically inject "status request" packets into the datastream (as shown at step 310 of FIG. 3), to which the subscribers need to respond. These packets may be piggybacked onto other packets relating to multicast publication messages. More accurately, the broker needs at least one subscriber to respond by sending a "status response". For maximum efficiency, it is preferable to minimise the number of subscribers who send status response packets in response to a status request packet.

15

If status packets are transmitted over the multicast fabric (i.e., are injected into the normal dataflow) then they are received (almost) simultaneously by all subscribers. To minimise the number of subscribers who respond to the status packet, the subscribers behave as follows:

20

Step 320 On receipt of a "status request" packet from the broker, each subscriber sets a short random duration backoff timer.

Step 330 On expiry of the timer a subscriber sends a multicast packet stating that it will respond to the broker, here called a "response claim" packet.

25

Step 340 If the subscriber receives a multicast response claim packet from another subscriber before the backoff timer expires, it cancels the timer and discards the status request packet and the other subscriber's response claim packet.

30

Step 350 A subscriber who has sent a response claim packet must establish a point to point connection with the broker and send a "status response" packet to the broker.

The broker may receive multiple status response packets, but the number should be minimised by the above algorithm.

35

Step 360 If a subscriber who has sent a response claim subsequently fails before managing to send the status response, or if it sends a status response but it doesn't get through to the broker, then the broker may receive no response to the status request. In this case the broker should re-send the status request before making a judgement about the state (i.e., the liveness) of the group of subscribers.

40

Technique #2

5 A second technique for liveness monitoring is similar to technique #1 described above, but with the addition of a mechanism to minimise the need for the broker to re-send status requests. This modification is based on the intent of a number of subscribers to respond, providing a degree of tolerance to subsequent subscriber failures.

10 The broker may optionally include a response quota in the "status request" packet, which includes a "number of replies (nr)" parameter (as shown at step 410 of FIG. 4, which will be described in more detail below). A subscriber with a pending backoff timer cancels it and discards the response only if it receives at least *nr* multicast response claim packets from other subscribers before the backoff timer expires. This will
15 guarantee that at least *nr* subscribers will try to respond to the broker, reducing the risk that the broker will have to re-send the status request.

20 Despite the intention of *nr* subscribers to respond, the broker will in general only need to handle less than *nr* incoming TCP connections or *nr* incoming UDP datagrams. This is because the broker, upon first successful status response reception, immediately sends a "response received" packet, over the same channel used to send the "status request" packet. Upon
25 "response received" packet reception, the subscribers with pending timers cancel them and discard their status packet.

25 The broker can escalate from technique #1 to technique #2 (with a response quota greater than 1) in the event of no responses being received within an acceptable time period.

30 As shown in FIG. 4:

Step 420 On receipt of a "status request" packet from the broker, that contains a "number of replies (nr)" parameter, each subscriber sets a short random backoff timer and initialises a response claim counter to zero.

35 Step 430 On expiry of the timer a subscriber sends a multicast packet stating that it will respond to the broker, called a "response claim" packet.

Step 440 If the subscriber receives a multicast response claim packet from another subscriber before the backoff timer expires, it discards the other subscriber's response claim packet and increments the response claim counter. If the counter reaches
40 *nr*, it cancels the timer and discards the status request packet.

Step 440 A subscriber who has sent a response claim packet sets a short backoff timer after which it will establish a point to point connection with the broker and send a "status response" packet to the broker.

5 Step 450 On receipt of a first status response by the broker, the broker sends a "response received" packet on the multicast channel.

Step 450 Any subscriber which is waiting for a backoff timer to expire can cancel the timer and discard the status packet.

If a subscriber who has sent a response claim subsequently fails before managing to send the status response, or if it sends a status response but it doesn't get through to the broker, then the broker should still receive a response from an alternative responding subscriber.

The broker may still receive multiple status response packets, but the number should be minimised (further) by the above algorithm.

Technique #3

20 A third technique provides a performance optimisation in the case where a TCP/IP connection is to be used for the subscriber-to-broker response channel. This technique can be used in combination with either of the techniques #1 and #2 described above.

25 During registration of a subscriber a TCP/IP connection is established between the subscriber and the broker. Once subscription (including key exchange, etc.) is complete the TCP/IP connection could be disconnected. This is beneficial for scalability. However, if at least some of the TCP/IP connections are maintained beyond the end of the subscription protocol, then they can be re-used for status response traffic, avoiding 30 the overhead of re-establishing a TCP/IP connection, which would be considerable (e.g., 7 packets to set up the connection compared to one status packet to be sent). Each TCP/IP connection can be associated with an idle timer and can be disconnected on expiry of the idle timer. Whenever a 35 connection is used (for subscription, key exchange or status traffic) the idle timer is reset.

40 Referring now to FIG. 5, as in technique #1 or technique #2 described above the broker sends a "status request" packet (step 510, which may contain a "number of replies (nr)" parameter). The inclusion of the nr parameter would require the same earlier-described counter logic, which for simplicity is not included in the description that follows. The subscriber behaviour is modified as follows:

- Step 520 On receipt of a "status request" packet from the broker, each subscriber checks whether it has an active connection to the broker. If it does, it either:
- A) immediately sends a multicast response claim packet, and sends a status response to the broker, or
 - B) sets a very short random backoff timer (shorter maximum delay than the minimum for backoff timers used by non-connected subscribers), and then sends a multicast response claim and status response.
- Option A) is simple but may result in more status response packets being sent to the broker; option B) should reduce this number of packets at the expense of some additional complexity in the subscriber code).
- Step 530 If a subscriber does not have an active connection to the broker it sets a short random backoff timer. On expiry of the timer a subscriber sends a multicast response claim packet. If the subscriber receives such a multicast response claim packet from another subscriber before the backoff timer expires, it cancels the timer and discards the status request packet and the other subscriber's response claim packet.
- Step 540 A subscriber that sent a multicast response claim packet establishes a point to point TCP/IP connection with the broker and sends a "status response" packet to the broker. This TCP/IP connection is left open for future status reports (from this moment on, this subscriber will know it has an active connection to the broker, and will respond more rapidly, according to the rules described above).
- The broker may receive multiple status response packets, but the number should be minimised by the above algorithm.
- Step 550 If a subscriber who has sent a response claim subsequently fails before managing to send the status response, or if it sends a status response but it doesn't get through to the broker, then the broker may receive no response to the status request. In this case the broker should re-send the status request before making a judgement about the liveness of the group of subscribers.

Technique #4

A fourth technique, alternative to technique #3 described above, contains a performance modification which is that the broker notes the identity of the first subscriber to register interest in a topic. The broker maintains the TCP/IP connection to this subscriber. Rather than

multicasting the status request packet, the broker sends it on the point to point TCP/IP connection directly to that designated subscriber.

5 If the designated subscriber fails then the broker will detect this because the TCP/IP connection will be broken. In this case the broker can revert to the multicast request scheme(s) above, and establish a new designated subscriber, as the first subscriber to respond to the status request.

10 It will be understood that in any of the above techniques it would be possible to use a custom reliable point to point protocol in place of UDP/IP or TCP/IP for the response channel from each subscriber to the broker.

15 It will also be understood that the broker may be arranged to be a listener in all multicast groups, so that it hears the 'claim' from subscribers, without any other explicit subscriber broker response being necessary.

20 It will be appreciated that the method described above for liveness monitoring in a publish/subscribe messaging system may be carried out in software running on a processor (not shown), and that the software may be provided as a computer program element carried on any suitable data carrier (also not shown) such as a magnetic or optical computer disc.

25 In summary, it will be understood that the techniques for efficient liveness monitoring in a reliable multicast system described above provides the advantage of improving the efficiency of network usage by reducing the number of unwanted packets that are sent.

CLAIMS

1. A publish/subscribe messaging system, comprising:

5 at least one broker and at least one subscriber, the broker having means for sending a status request message to the subscriber, and means for receiving an indication of liveness of the subscriber.

10 2. The system of claim 1 wherein the subscriber has means for sending a status response message to the broker to indicate liveness.

15 3. The system of claim 1 wherein the means for receiving comprises means for listening on a multicast channel and for determining an indication of non-liveness from failure to receive a response from the subscriber.

20 4. The publish/subscribe messaging system of claim 2 wherein the means for sending a status response message to the broker comprises means for suppressing sending of the status response message if at least another subscriber sends a status response message.

25 5. The publish/subscribe messaging system of claim 4 wherein the means for suppressing sending of the status response message comprises:

means for setting a timer upon receipt of a status request message from the broker;

25 means for sending, on expiry of the timer, a multicast message claiming response to the broker;

30 means for cancelling the timer and discarding the status request message if the subscriber receives a message claiming response from another subscriber;

35 means for sending the status response message to the broker following sending a message claiming response.

35 6. The publish/subscribe messaging system of claim 5 wherein the broker further comprises means for re-sending the status request message if it does not receive a response thereto.

40 7. The publish/subscribe messaging system of claim 4 wherein the means for suppressing sending of the status response message is arranged to suppress sending of the status response message if at least a desired plurality of other subscribers send a status response message.

8. The publish/subscribe messaging system of claim 7 wherein the means for suppressing sending of the status response message comprises:

5 means for setting a timer upon receipt of a status request message from the broker, the status request message containing a parameter representative of the desired plurality of other subscribers;

10 means for sending, on expiry of the timer, a multicast message claiming response to the broker;

15 means for cancelling the timer and discarding the status request message if the subscriber receives messages claiming response from the desired plurality of other subscribers;

20 means for sending the status response message to the broker following sending a message claiming response..

25 9. The publish/subscribe messaging system of claim 5 or claim 8 wherein the timer has a random duration.

10. The publish/subscribe messaging system of any one of claims 1-9 wherein at least one of the subscribers is arranged to maintain an active connection to the broker established during registration, and to use the active connection to indicate liveness to the broker.

30 25 11. The publish/subscribe messaging system of claim 10 when dependent on claim 4 wherein the means for suppressing sending of the status response message comprises:

35 30 means for checking, upon receipt of a status request message from the broker, whether the subscriber has an active connection to the broker and if so performing one of A) and B):

A) sending a multicast response claim message, and sending a status response message to the broker via the active connection, and

B) setting a timer and then sending a multicast response claim and a status response message to the broker via the active connection; and

40 means for, following sending of a multicast response message, establishing an active connection to the broker if not already established and sending the status response message to the broker via the active connection.

12. A publish/subscribe messaging system according to any preceding claim, wherein the broker is arranged to designate as a primary subscriber the first subscriber to register interest in a topic, and to maintain an active connection to the primary subscriber for sending directly to the primary subscriber a status request message, and in the event of failure of the primary subscriber to send a status request message to at least one other subscriber and to designate as a new primary subscriber the at least one of the other subscribers whose indication of liveness is next first received.

5

10 13. The publish/subscribe messaging system of claim 10, 11 or 12 wherein the active connection is a TCP/IP connection.

15

14. The publish/subscribe messaging system of any preceding claim wherein the status request message is piggybacked onto another multicast publication message.

20

15. The publish/subscribe messaging system of any preceding claim wherein the indication of liveness is sent over one of:

- a UDP connection, and
- a TCP connection.

25

16. The publish/subscribe messaging system of claim 15 wherein the connection over which the indication of liveness is sent is arranged to escalate autonomously from a UDP connection to a TCP connection in the event of no responses being received by the broker within a chosen time period.

30

17. A method for liveness monitoring in a publish/subscribe messaging system having at least one broker and at least one subscriber, the method comprising:

35

at the broker, sending a status request message to the subscriber, and receiving an indication of liveness of the subscriber.

18. The method of claim 17 comprising:

40

at the subscriber, sending a status response message to the broker to indicate liveness.

19. The method of claim 17 wherein the step of receiving comprises listening on a multicast channel and determining an indication of non-liveness from failure to receive a response from the subscriber.

5 20. The method of claim 18 wherein the step of sending a status response message to the broker comprises suppressing sending of the status response message if at least another subscriber sends a status response message.

10 21. The method of claim 20 wherein the step of suppressing sending of the status response message comprises:

setting a timer upon receipt of a status request message from the broker;

15 one of A) and B):

A) sending, on expiry of the timer, a multicast message claiming response to the broker;

20 B) cancelling the timer and discarding the status request message if the subscriber receives a message claiming response from another subscriber;

25 sending the status response message to the broker following sending a message claiming response.

22. The method of claim 21 further comprising, at the broker, re-sending the status request message if the broker does not receive a response thereto.

30 23. The method of claim 20 wherein the step of suppressing sending of the status response message comprises suppressing sending of the status response message if at least a desired plurality of other subscribers send a status response message.

35 24. The method of claim 23 wherein the step of suppressing sending of the status response message comprises:

40 setting a timer upon receipt of a status request message from the broker, the status request message containing a parameter representative of the desired plurality of other subscribers;

one of A) and B):

A) sending, on expiry of the timer, a multicast message claiming response to the broker;

5 B) cancelling the timer and discarding the status request message if the subscriber receives messages claiming response from the desired plurality of other subscribers;

10 sending the status response message to the broker following sending a message claiming response.

10 25. The method of claim 21 or claim 24 wherein the timer has a random duration.

15 26. The method of any one of claims 17-25 further comprising, at least one of the subscribers, maintaining an active connection to the broker established during registration, and using the active connection to indicate liveness to the broker.

20 27. The method of claim 26 when dependent on claim 20 wherein the step of suppressing sending of the status response message comprises:

25 checking, upon receipt of a status request message from the broker, whether the subscriber has an active connection to the broker and if so performing one of A) and B):

25 A) sending a multicast response claim message, and sending a status response message to the broker via the active connection, and

30 B) setting a timer and then sending a multicast response claim and a status response message to the broker via the active connection; and

35 following sending of a multicast response message, establishing an active connection to the broker if not already established and sending the status response message to the broker via the active connection.

28. A method of liveness monitoring in a publish/subscribe messaging system having at least one broker and at least one subscriber, comprising:

40 at the broker, designating as a primary subscriber the first subscriber to register interest in a topic, and maintaining an active connection to the primary subscriber for sending directly to the primary subscriber a status request message, and in the event of failure of the primary subscriber to revert to the request method of any one of claims

17-27 and designating as a new primary subscriber the subscriber whose indication of liveness is next first received.

29. The method of claim 26, 27 or 28 wherein the active connection is a
5 TCP/IP connection.

30. The method of any one of claims 17-29 wherein the status request message is piggybacked onto another multicast publication message.

10 31. The method of any one of claims 17-30 wherein the indication of liveness is sent over one of:

a UDP connection, and

15 a TCP connection.

32. The method of claim 31 wherein the connection over which the indication of liveness is sent escalates autonomously from a UDP connection to a TCP connection in the event of no responses being received by the
20 broker within a chosen time period.


33. A computer program element comprising computer program means for performing substantially the method of any one of claims 17-32, when said program is run on a computer.


25 34. A publish/subscribe messaging system substantially as hereinbefore described with reference to any one of FIGS. 3-5 of the accompanying drawings.


30 35. A method, for liveness monitoring in a publish/subscribe messaging system, substantially as hereinbefore described with reference to any one of FIGS. 3-5 of the accompanying drawings.


ABSTRACT**LIVENESS MONITORING IN A PUBLISH/SUBSCRIBE SYSTEM**

5 A variety of techniques are disclosed for efficient liveness monitoring in a reliable publish/subscribe multicast system having at least one broker and at least one subscriber, by: at the broker, sending (310) a status request message to the subscriber, and at the subscriber, sending (350) a status response message to the broker to indicate liveness. Status responses by subscribers may be suppressed when at least a chosen minimum number ($>=1$) of subscribers send a response. This suppression (320-340) uses a system of "response claim" messages and random back-off timers. These provide the advantage that the efficiency of network usage is improved by reducing the number of unwanted packets that are sent.

**FIG. 1**

**FIG. 2**

**FIG. 3**

**FIG. 4**




FIG. 5